

BullSequana X series

BullSequana X410 E5 dense GPU-accelerated compute node



The BullSequana X400 series is a comprehensive family of rack-mounted servers designed for High Performance Computing and offering an optimum balance between cost and efficiency.

Derived from technologies initially developed for game machines, **NVIDIA® Tesla® GPUs** (graphics processing unit) are highly parallel processors designed to boost computing.

GPU-accelerated computing is the use of a GPU together with a host CPU to accelerate High Performance Computing, deep learning, and analytics applications. Pioneered in 2008 by Bull in its supercomputing range, GPU accelerators now power energy-efficient data centers in government labs, universities, enterprises, and small-and-medium businesses around the world. They play a huge role in accelerating applications in platforms ranging from HPC and artificial intelligence to cars, drones, and robots.

GPU-accelerated computing offloads compute-intensive portions of the application to the GPU, while the remainder of the code still runs on the CPU. From a user's perspective, applications simply run much faster.

Implementing GPU computing is getting easier every month as new tools become available, but it is still not trivial. To help you innovate with GPUs with maximum payoff, Atos offers complete GPU-based solutions:

- hardware leveraging the latest technologies,
- an integrated software environment,
- expertise and services to implement your solution.

Maximized density for maximized performance

The new 1U BullSequana X410 E5 servers are optimized for compute density, performance and power efficiency:

- They support up to four NVIDIA® Tesla® P100 GPUs, interconnected either via PCIe direct connect or via NVIDIA® NVLink™.
- It also features two future Intel® Xeon® Scalable processors.
- It is powered by ultra-high efficiency redundant Titanium Level power supplies.

The integration of up to four GPUs within 1U places this system at the forefront of today's accelerated computing solutions, both in terms of performance and cost-optimization.

High Performance GPU communication

The BullSequana X410 E5 servers are available in two versions:

- With traditional PCIe: a streamlined layout architecture enables PCIe direct connect for best signal integrity as well as elimination of complex cabling.
- With NVIDIA® NVLink™ interconnect: NVLink brings performance benefit for GPU-to-GPU communications. This technology increases exponentially application performance: NVIDIA® announces that it allows data sharing at rates 5 to 12 times faster than the

traditional PCIe Gen3 interconnect.

Seamless integration in the HPC environment

The Bull SuperComputer Suite integrates all the tools needed to operate a hybrid system comprising GPUs and CPUs. The task management environment supports the allocation of applications to the relevant computing resource. An application can naturally use both computing resources so as to fully capitalize on the rich potential offered by a hybrid system.

Expertise & services for optimal performance acceleration

Depending on the type of application and its degree of optimization, GPUs can accelerate processing by a factor of 1 to 100! Expertise is essential to get the most out of GPUs. Atos has a long experience of systems associating GPUs and CPUs: our engineers have had hands-on experience of porting and optimizing applications for accelerated computing for many customer projects.

Atos' **Center for Excellence in Parallel Programming** delivers the highest level of expertise and skills to help organizations optimize their applications to take advantage of GPUs. It provides a large portfolio of services, from analysis, consulting, parallelization, to optimization of application codes.

Technical specifications

BullSequana X410 E5 with PCIe GPUs



BullSequana X410 E5 with NVLink GPUs



Form Factor	Rack-mounted 1U drawer	
Compatible GPUs	All cards must be identical - no mix of technologies: <ul style="list-style-type: none"> • up to 4 Nvidia® Tesla® P100 GPUs (dual slot) OR • up to 4 Nvidia® Tesla® P40 GPUs (dual slot) 	
Processors	2 Intel® Xeon® processors Scalable family	
Architecture	Future Intel® chipset Future Intel® interconnect	
Memory	Up to 16 DIMM slots (8 DIMMs per CPU) supporting latest generation ECC DDR4 DIMM modules, and ready for future memory technology Support of the highest available memory transfer rate	
Expansion Slots	4 internal PCI-E 3.0 x16 for GPU + 2 PCI-E 3.0 x8 in x16 slot	2 PCI-E 3.0 x16 for RDMA + 1 PCI-E LP x16 +1 PCI-E LP x8
Storage Devices	2 front-access slots for: <ul style="list-style-type: none"> • 2.5" SATA3 hot swap hard disk drives OR SSD Flash Disks 2 internal- slots for: <ul style="list-style-type: none"> • 2.5" SATA3 hard disk drives OR SSD Flash Disks OR 2.5" internal NVMe drives <ul style="list-style-type: none"> • 1x M.2 NVME drive 	
Network	Dual Port 10G BASE-T LAN Optional: up to 2 InfiniBand FDR or EDR / OPA PCI-E, single/dual port adapter cards (PCI-E x8)	Dual Port 10G BASE-T LAN Optional: up to 2 InfiniBand FDR or EDR / OPA PCI-E, single/dual port adapter cards (PCI-E x16)
Security	Chassis intrusion detection	
I/O Ports	Back Panel connectors and I/O ports: <ul style="list-style-type: none"> • 2 USB 3.0 ports • 2 Gigabit Ethernet RJ45 ports • 1 IPMI/RedFish RJ45 port 1 VGA port	
Management	1 BMC (IPMI 2.0 with virtual media-over-LAN) - Embedded Aspeed 2500 with dedicated and LAN shared RJ45 ports	
Power Supply	2 x 2000W 80PLUS Titanium Level redundant PSU	
Ventilation	9x counter-rotating fans with optimal fan speed control	
Physical Specifications	Dimensions: 43mm x 437mm x 778 mm (H x W x D) Gross Weight: 21.8 kg (maximum)	
OS & Cluster Software	Red Hat Enterprise Linux & Bull SuperComputer Suite	
Regulatory Compliance	Safety (EC, IEC, UL, CSA certifications) ; Electromagnetic Compatibility (EC, FCC, ICES-03, VCCI certifications), Environment (RoHS II & WEEE directives, REACH regulation)	
Warranty & Services	Standard server warranty 1 year, extended: consult us	

Find out more about us
hpc@atos.net
atos.net/products/high-performance-computing-hpc

© Atos October 2017 - All trademarks are the property of their respective owners. Atos, the Atos logo, Atos Codex, Atos Consulting, Atos Worldgrid, Bull, Canopy, equensWorldline, Unify, Worldline and Zero Email are registered trademarks of the Atos group. Atos reserves the right to modify this document at any time without notice. Some offerings or parts of offerings described in this document may not be available locally. Please contact your local Atos office for information regarding the offerings available in your country. This document does not represent a contractual commitment.